Release 1.0

86-4 21 April 1986

SEGMENTING CEPS

We recently attended the corporate electronic publishing systems (CEPS) revival meeting in Los Angeles. Despite the huge brouhaha about this new market (the salvation of Apple, the next mass-market/horizontal application, etc.), the show attracted only a few thousand people and seemed remarkably tame. With minor exceptions, attention focused on the mingling of text and graphics rather than on attempts to control that process. Rather than ask the practical questions of someone familiar with the publishing process (How will this help my bottom line? How many users does it support? How will it fit in with what I've got?), customers, vendors and press are focusing on the gee-whiz: "Merge text and graphics at the touch of a button!" "Supports 300 x 300!" "Make every employee a graphic artist!"

Unlike most healthy markets we've seen, CEPS has had no chance to reach maturity in the forgiving obscurity of niche-hood. Rather, it's like a garden of Easter eggs laid out for the kids...now overrun with grown-ups -- venture capitalists, analysts, competitors, Xerox, DEC, Atex/Kodak, and the like. Companies who don't raise money, form alliances and find distribution channels will be left out. That's why Ventura signed its deal with Xerox, Interleaf took $3.5 million from Kodak, and MicroPro has well-founded hopes of becoming parent to one or other of the smaller page-layout package vendors.

Mixing markets

In truth, CEPS is not a market so much as a function *within* a market -- the corporation. (Newsletter publishers are a market; so are newspaper publishers, but they have totally different needs.) In the corporate market, there's *personal* publishing, and there's *corporate* publishing. Personal publishing lets users create pretty documents, frequently with text and graphics merged, and print them out for all to see. Personal publishing helps people to make more effective presentations, internal and external, and to insult their friends by sending them personal letters that look mass-produced. Personal publishing can be accomplished

⟹

THIS SPACE FOR RENT

INSIDE

SEGMENTING CEPS 1
Easter eggs. thINKer, thINKertools.
BUSINESSLAND LANDS MBI 7
Next, Businessplanet?
IBM'S CONVERTIBLE 8
...and some other items.
INTERSTITIAL PROGRAMMING 9
Finding work for idle cpus.
SYMBOLICS DELIVERS A VEHICLE 10
The Applications Engine.
RETROACTIVE PRODUCTIVITY 12
Combing spaghetti.
RELEASE 2.0 14
Silogic restructured.
RELEASE 0.5 16
SmartNotes, Smart maps, Squeeze me, iLINK me, Programmed PR, TI ties, IntelliShuffle, Xerox AI.
CALENDAR; PHONES 21/23

Ziff-Davis Publishing Company, One Park Avenue, New York, NY 10016, (212) 503-5500

nicely on a Macintosh, and even cost-effectively if several users are hooked up to an Apple LaserWriter. Personal publishing is like pcs themselves; it can creep into a company disguised as office equipment on a purchase order; it can be espoused by one person and lure others; it can on occasion sharply increase individuals' productivity. Although personal publishing is frequently used to improve the form of the same old content, it may help to make that content more intelligible or at least more interesting to others.

Corporate electronic publishing is another matter entirely. It concerns not pictures and texts, but copyflow management, project management, version management, edit cycles, indexing, etc. It's used to produce not letters, newsletters, price lists and sales memos, but manuals, formal contract proposals, catalogues, sales literature, journals and the like. These systems should be multi-user (whether peer-networked or hosted), purchased with the consent of top management. They produce products that generate revenue, and the production process is closer to a factory's than to a jobshop's.

But the current set of single-user text/graphics editors, page-layout systems and the like don't really do the job, even when attached to a mutli-user data base. We need the sort of systems long familiar to newspapers, which control copy flow, deadlines, version management, length, and the like. These systems exist at the high end, from companies such as (Kodak's) Atex, Digital Equipment and the like; they could also be built around PCs and the MacPlus and its successors -- if only someone would do the software. Should 3Com buy Aldus?

Indeed, we have spent much of the last ten years on "office automation" and "the human interface." But something fundamental is missing. We have automated clerks and clerical tasks, not the complex interactions among people and processes that actually go on within an office. It's now relatively easy for a person to talk to a computer. But what about one computer/application talking to another? This is what Lotus, and others, are working on under the name of group-productivity software -- and what the CEPS vendors need to provide to their customers.

Leading the charge

Electronic publishing is an ideal lab for testing the proposition that dealers can successfully move upmarket, and it may be the "application" that finally boosts local-area network sales. Electronic mail becomes addictive once people get exposed to it, but there is generally no compelling reason to start using it (without a committed proponent somewhere in an organization). Electronic publishing is different: It's not just a capability, but an application. It's a job that has to get done -- and often requires the resources of several people, communicating, to be accomplished.

Who will succeed in this market? We expect to see the word-processing companies -- Multimate/Ashton-Tate, Microsoft, Samna, Satellite Software -- come out with products that do 80 percent of the work with 20 percent of the cost and hassle. The IBM PC, perhaps not the ideal system for this exercise, will attract software to perform 80 percent of what the Mac can do, for 20 percent of the hassle (if not 20 percent of the cost). Since a benefit of desktop publishing is the ability to integrate discrete documents and formats, it makes sense to put it on the PC where all your other data is available. Ultimately, the difference between the high and the low end will reside more in software, support and output devices than in the computers.

HELLO SWEETHEART! GET ME THE VT-241!

Addressing the issue of multi-system interaction at the very highest end of the publishing market is Composition Systems, Inc., a longtime vendor of front-end* systems to the newspaper community, with revenues of $20 million last year. CSI is using Inference Corp.'s Automated Reasoning Tool (ART) to move beyond its traditional line and automate the management as well as the operations of its customers. Its system will cost somewhere between $1 and $3 million depending on configuration, and naturally enough will sell only to customers for whom publishing is a major business (not a support function) and who must run their operations and lay out their product in real time -- in other words, newspapers. (A catalogue publisher has the luxury of planning ahead.) As this "expert publishing project" as yet has no name, we have unilaterally nicknamed it "thINKer" for the purposes of this piece.

----+-----

The current vision of AI is of a guru sitting in a corner, a helpless machine that requires continual stimulation by a user ("interactive consultation") to do useful work. In fact, expert systems should spend most of their efforts querying production data bases and guiding existing applications to do their work more effectively, efficiently, and flexibly. How do we integrate such systems into the dp world?

-----+----

CSI's project, notable for its ambition, is also notable for its structure (see page 4). While thINKer is intensely particularized to the operations of a newspaper business, it addresses many problems and issues of more general interest, including how to get expert systems to interact with (stupid?) production applications. Indeed, CSI is now selling the underlying tools that it built to do just that for this VAX-based system (see page 6). Project director Mike Stock knows that a newspaper is not just a factory with copy going in at one end and papers out the other, but rather a network of interacting units with conflicting goals, data, constraints, etc.

thINKer comprises three modules -- layout manager, production control, and space reservations -- each of which assumes the roles of several decision-makers within a newspaper: the production manager, the space manager, the editor, the edition manager, and the like. The ad salesman has to contend with the edition manager, who knows that if he adds a single ad to the Monday morning north suburban edition, he'll have to fill two extra pages. He suggests holding the ad until Wednesday... The dummying manager keeps the alcohol ad away from the story about the five-car pile-up after a high-school drinking spree... The production manager readies the presses to handle five editions, including a color insert about Bermuda... The editorial module queries the writers to see if their copy will be in on time, and the right length... The layout module and the production control module together rejigger the paper and presses to accommodate the deadline-busting scoop about a local airline takeover engineered by First Boston...

**Front-ends are terminals with which operators type classified ads, typesetters type editorial copy, and, most recently, commercial artists build ads and editorial illustrations. One winner in this last category is Boston's Camex, which has sold a system to the New York Times and should sell one to 47th Street Photo, which still creates its ads by hand.*

thINKer High-Level Control Flow Diagram

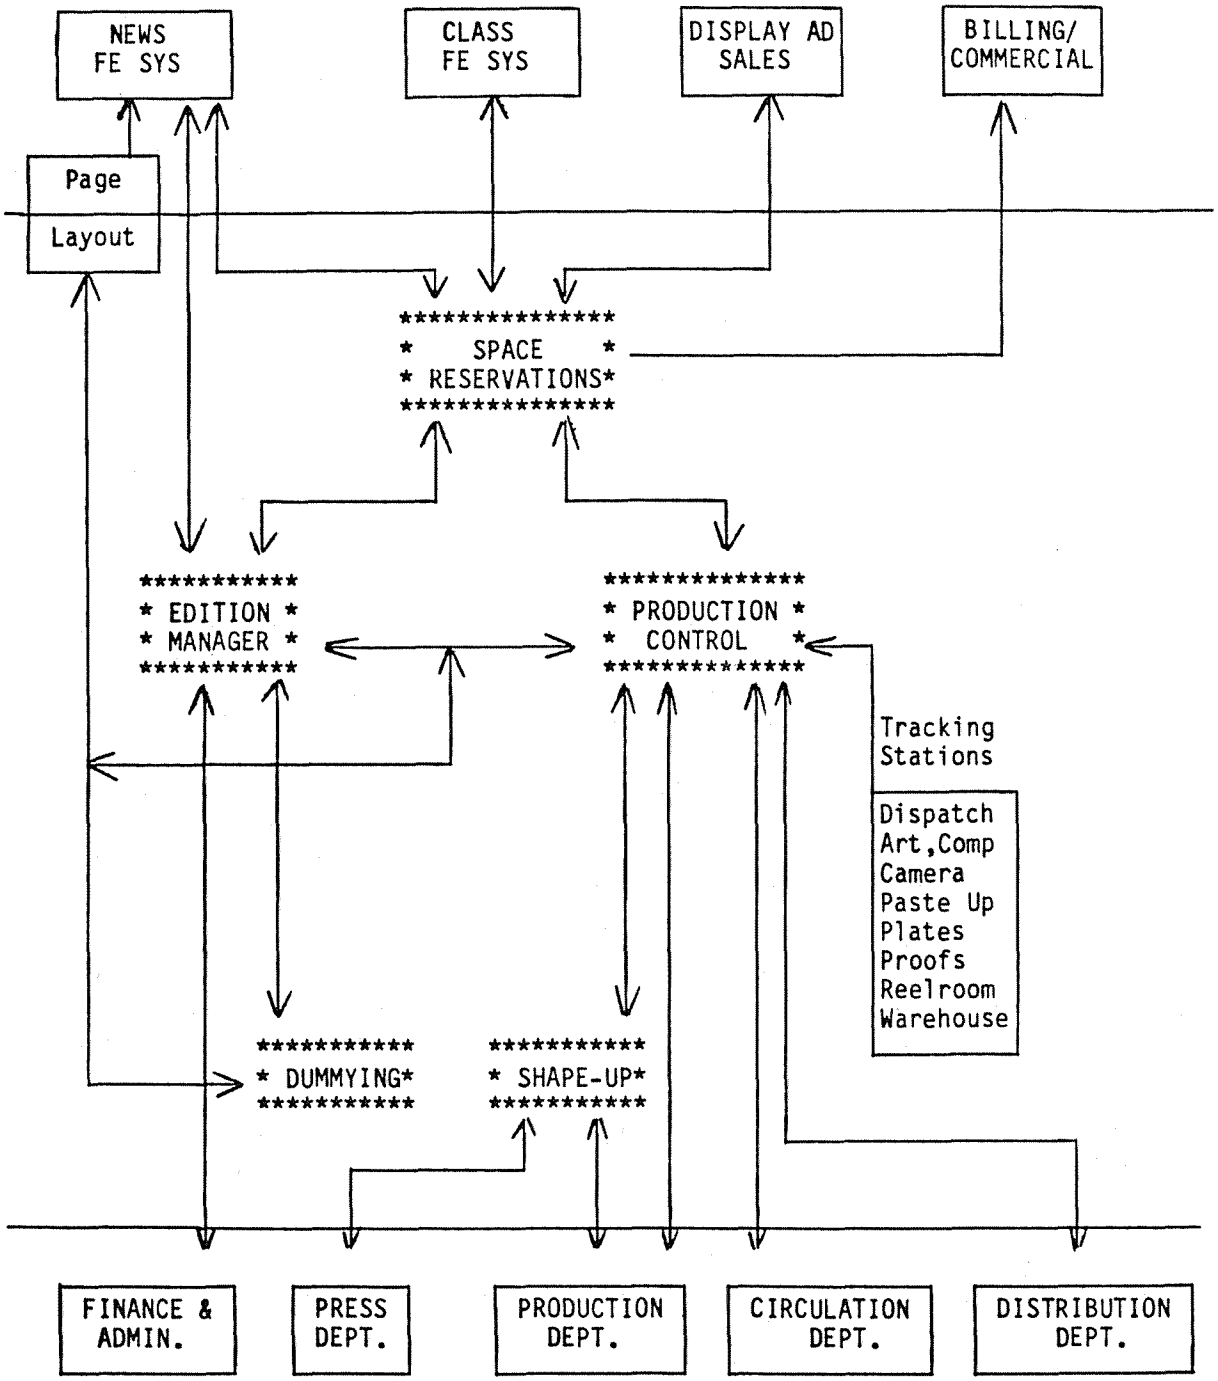

"Experts" are shown surrounded by stars; operating and production departments (people and applications) are in boxes.

Source: Composition Systems, Inc.

From a practical point of view, this modular implementation makes sense: Modules can be installed and integrated into a customer's operations one at a time, with minimum disruption. Ultimate tangible benefits include a higher proportion of paid ads (versus make-goods because some element of timing or placement was mishandled), greater flexibility in changing ads or editorial at the last moment, and improved ability to create zoned editions. Jobs may ultimately be eliminated, less ink and paper wasted, use of color optimized, and so forth. Most of all, thINKer should improve working conditions and product quality: As CSI's sales literature so delicately states, thINKer "evaluates the implications of unscheduled events...to eliminate bottlenecks, last-minute heavy work volume, and the inflated staffing presently needed to cope with these situations."

Inside thINKer

thINKer, built on a foundation of Inference's ART, runs on VAXes with DEC's VAX LISP. VAXes are both easier to sell into a production environment and higher-performance than LISP machines, contends Stock, a reformed academic. (We'll concede the first, at least.) But there's no question of using installed equipment; thINKer will load a VAX fully, with no capacity left over. In terms of lines of code, the applications are roughly 60 percent ART, 30 percent LISP, and 10 percent other, while the underlying structure is about 60 percent LISP, 30 percent other, and 10 percent ART. In terms of resources used, the expert system component does most of the work, applying its rules to real-world data.

In each module there's a representation of that module's "world:" costs, page real estate, dates and times of editions, ads and editorial plans. The space manager, for example, maintains a list of advertisements, with their sizes, color requirements, whether an advertiser is "preferred," requested placement (business section, inside front page, etc.), timing (Wednesday's food section, or a certain week for a time-sensitive promotion), the charges (linked into a billing system), etc. Classifieds are grouped into categories and treated as a mass by the space system, which doesn't need to know details of interest only to the invoicing and accounting departments. Cancellations and additions, changes and queries from users and from the other expert modules continually bombard each world with new data, and even with information leading to the formation or modification of rules.

Most of the information is stored in plain old relational tables managed by DEC's RDB dbms. Needed data are searched out by an expert module and instantiated into the knowledge base to be acted on by the rules: How are priorities assigned? How are pages laid out? What are the constraints governing each decision?

An ad salesman whose client has gone bankrupt can type in, "Kill the Computer Confusion ad!" and watch on a DEC VT-241 graphics terminal as thINKer places ads on the pages and shuffles them around to arrive at the optimal arrangement -- satisfying as many requests as possible and meeting all the constraints. (It also knows how to fill the extra half-inch with "Remember the Neediest!" or "Subscribe to Computer Industry Daily!")

This sounds simple enough as described; the trick is to manipulate all this application-specific knowledge -- rules and relationships, data base templates, the lexicon, etc. -- in real-time, meeting real constraints. To

achieve this, CSI's underlying toolkit links the various modules together, and provides a link to humans too, via natural-language translation. Which module takes precedence? That's programmed into the Meta module (part of the toolkit), which functions as the publisher, reconciling the conflicting claims of all the paper's constituents.

Commercial status

Currently, thINKer is under development at the "Bunker," a shabby office building near CSI's headquarters in Elmsford (Tarrytown), NY. Mike Stock has been working on it with a team of 17 programmers and newspaper experts for two years. The programmers program; the experts tell them how a newspaper works and continue to spend much of their time at customer sites. thINKer is designed so that customer "experts," newspaper personnel who want the system to reflect their own preferences, priorities and practices, will be able to talk to it directly through a combination of natural language, forms fill-in, and menu responses.

Customers for CSI's traditional equipment include the Wall Street Journal, the Newark Star-Ledger and the Peninsula Times-Tribune, but the company only recently disclosed its first live prospect for thINKer -- Southern Connecticut Newspapers, Inc., a Times-Mirror subsidiary and publisher of the Stamford Advocate and Greenwich Time. SCNI and CSI together are working on a feasibility demonstration at the Bunker. If all goes well, shipments should start late this summer to SCNI and other papers. Modules will cost about $80,000 each, hardware extra. A three-module system will set a typical 50,000-circulation paper back a million or so, but Stock is confident in saying the system will pay for itself within three years.

Will all go well? It's hard to believe all will go *smoothly*, but the basics of the system seem to be in place. The demo we saw was still excruciatingly slow, but CSI is trying to prove it can be done, elegantly or not. In fact, the minute or two the system takes to reconfigure the entire paper is a miracle when you consider how -- and how ineffectively -- *people* do it.

thINKertools

thINKer's toolkit incorporates a significant sunk investment in systems expertise to get its numerous components to cooperate and to handle standard-format data, natural language, etc. Prices range from $6000 for a low-end workstation to $24,000 for a VAX 8600. Customers include DEC and three other companies who want to keep their interest quiet, Stock says.

We're not yet convinced of CSI's ability to sell its tools effectively to other systems designers once the market broadens, but we haven't yet seen the results of the company's current efforts to set up and fund a means of doing so. DEC, for one, has expressed interest in marketing the toolkit -- and would be a natural for the job.

Release 1.0 is published at least 15 times a year by EDventure Holdings, 375 Park Avenue, New York, NY 10152. (Note new address, effective June 1.) Editor & publisher: Esther Dyson; associate publisher: Sylvia Franklin.
Copyright 1986, EDventure Holdings Inc. All rights reserved. No material in this publication may be reproduced in any form without permission.

BUSINESSLAND LANDS MBI

Businessland's pending purchase of MBI, following hard upon its purchase of United Telecommunications' Midwestern AmeriSource chain, is simply the realization of a fundamental urge to expand. While AmeriSource is "broke" in the sense that it needs some fixing, MBI is a much cleaner acquisition, likely to require less management time from a company that will surely be stretched by the absorption of these two chains. (Moreover, MBI itself just recently completed the acquisition of Computer South, a chain of 11 stores.)

MBI gives Businessland the kind of presence in upscale Eastern markets (except New York City) that you can now only *buy*: Current conditions of overstoring and new-authorization moratoriums make it impossible to grow such a presence from scratch anymore. MBI also has a strong presence in corporate accounts and in the Federal government. And courtesy of its acquisition of Microsource (last year, it also has significant expertise in CAD/CAM, one of the vertical-market businesses so alluring to Businessland and its competitors.

In 1983 Businessland acquired just three stores in Phoenix in an expansion that proved difficult for the company to absorb (Release 1.0, 16 February 1986). Businessland has always had a strong corporate culture but was then too small to impose it easily on a foreign body. Now the company is far bigger. But is it capable of ingesting 72 outlets, for a total of 141 nationwide?

The upheaval may be as great as that wreaked by the three stores in 1983 -- troubling, but not permanently damaging. Businessland still lacks a president since the recent departure (following the slightly less recent arrival) of ex-IBMer Ron Watkins, in search of an organization where he'll have more control. Given Businessland chairman Dave Norman's record so far, it's hard to argue he should delegate more, but revenues of $600 million are enough to tax anyone's energies.

Market dynamics

Certainly, Businessland's impetus to grow is appropriate. At a time of consolidation, when market share is easily purchased, those who pick it up are likely to emerge the long-term winners.

For most retailers -- or local-sales companies for those who eschew the other term -- economies of scale are primarily market-based. That is, 30 percent of three markets is worth immeasurably more (probably infinitely more) than 3 percent of 30 markets -- a spread-out condition likely to produce nothing but losses. Economies of advertising, distribution, sales-force deployment, customer access to service and training, all come into play. But if you're operating on a national scale such as Businessland's, wooing Fortune 500 companies in competition with IBM itself, it helps to be able to say, "...and of course our Washington office can handle your lobbyists, our New York office can handle your advertising team, our Cupertino office can handle your R&D facility, and our Palm Beach office can handle your retired chairman's needs."

In other words, if you can afford it, 30 percent of 30 markets is the best of all.

IBM'S CONVERTIBLE

IBM finally announced its laptop convertible, with no surprises except for the date (April 2, a Wednesday, the day after April Fools' Day). It seems like a nice system -- we would gladly use one on loan -- but we find it hard to believe that even IBM's blessing can create this market.

The big issue here is the emergence of the 3.5-inch disk as a standard -- allowing us all to eschew portable computers in favor of carrying around a couple of virtually indestructible (hard) floppies to load into what ever computer we find (PC or compatible, to be sure) at our destination.

Compatibility overrated?

Second point: IBM intends to make life difficult for the clones. IBM alone has the market sway to get all the significant software vendors to rejigger their software for its new systems. (IBM will also use the weapon of price -- or, to put it politely, its advantages in low-cost manufacturing, as displayed in tapes of the company's highly automated production facilities in Austin.)

Indeed, absolute compatibility is overrated -- at least for IBM and friends. Think about it... IBM and the larger software vendors have no problem with customers having to buy new versions of old software -- even though the customers are not overjoyed. Most large software vendors will offer registered customers some kind of a trade-in or discount for 3.5-inch software; it's too good a marketing opportunity to pass up. Computer Associates, for one, is charging registered owners $50 for 3.5-inch versions of their software; Borland is charging a Borland-like $10; Microsoft is offering free trade-ins -- in return for an updated mailing list, that is.

Long-run, with the advent of the (mostly compatible) 386 and especially the RT PC, the segmentation of the business pc market into high and low ends will become unmistakable. IBM's recent announcement of a version of TopView that swaps excess data and programs to disk, recent improvements in software to aid connectivity between its PCs and IBM hosts, and continual increases in basic system configurations, encourage and foreshadow this trend.

At some point, it becomes foolish to try to make every piece of software work on every machine or congifuration. Unless you are, say, Lotus, you can't sell to everyone anyway, so why not design a product to take advantage of the configuration typical of your target customer? Javelin, Ansa, and a host of others have done so; Symantec addresses a less sophisticated market but provides an add-on option for memory-short customers. The message: It is no longer always practical or desirable to address the entire range of users with a single product -- hardware or software. That news, of course, is best for for very small or very large (guess who?) companies.

In the case of the Convertible in particular, and conspiracy theories to the contrary, everything seems to be fairly compatible with other PCs, clones, and portables. GRiD's software, for example, runs quite nicely on the Convertible, GRiD says. And software generally requires only reformatting, not recoding, to run on the Convertible. There are some rumblings about the system's proprietary ROM BIOS, but we can't yet find any situations where the differences matter.

INTERSTITIAL PROGRAMMING

We attend a lot of conferences and hear a lot of predictable speeches: Rod Canion on standards (compatibility being passe), John Sculley on innovation and connectivity, IBM on the importance of the customer, Randy Davis and Pat Winston on expert systems, and lots of folks on the year (or the millennium) of the network, information as a resource, and the like. So it was a pleasure at the Lotus Developer Conference last month to be surprised at the answer when we asked Jerry Kaplan, "Same old AI speech?" and he said "No."

Kaplan's pitch is not a new one: We've heard the same thoughts from Bill Gates among others, but more as a throwaway among a list of neat ideas (yes, you *can* be too prolific, at least in a single speech). Kaplan, a co-founder of Teknowledge and now Principal Technologist at Lotus, presented his concept with the proper positioning and, better yet, to the proper audience -- some 450 independent software vendors and internal corporate developers. In essence, he's propounding what we term "interstitial programming" -- the idea that personal computer processors spend most of their time waiting idly for a user to punch a key (which takes about a millisecond of the computer's attention), so for goodness' sakes let's give those processors something to do. ("An idle cpu is the devil's playground," Kaplan warned.)

Certain terminate-and-stay-resident programs that pop up as appropriate are one approach to that issue, though Kaplan didn't mention that specifically. (So is incremental garbage collection and so is the use of pipelining, at a much lower level, in sequencing the execution of instructions by a microprocessor.) Within an application there are tremendous wasted resources that could be devoted to such useful tasks as anticipating what screen the user might like next (the next chunk of text in a word-processing program, for example), pre-sorting a column of numbers, or the like. Obviously, the program has to know a little more about how a user and, better yet, a specific user, is likely to use the program.

Microsoft already makes use of these techniques: Word (Mac and PC versions) gives the impression of extremely rapid scrolling simply by having the surrounding screenfuls at the ready. To the extent it can, Multiplan first calculates the cells that are on display and does the rest later in private.

With a little bit of pattern recognition and the disk space to store data from session to session, these capabilities are not hard to build in. The greater effort will be in writing applications with a view to being "interruptible," notes Kaplan, and also paying attention to how you can do parts of a task so as to give the appearance of completion. For example, a sort normally isn't "soup" until it's done. But suppose instead of doing a normal sort the system first selected all the "A" words, and sorted only them -- a much quicker task. While the system was displaying the sorted "A" words, it would go on to sort the rest of the file in background, giving the illusion of an extremely fast sort. (If the user has a tendency to flip through his data quickly, the system could go on to the "B" words, and so forth. It takes longer to do the full sort this way, but the user never notices because he is looking at the results incrementally all the while.)

The perceived performance improvements of interstitial programming are enormous. Moreover, a program "tuned" for its regular owner would be less satisfying to anyone other than that owner. Interesting implications there...

JOINING THE FRAY: SYMBOLICS DELIVERS AN APPLICATION VEHICLE

In the good old days, AI was easily defined as whatever ran on a LISP machine. As the leader in the symbolic-computer marketplace, with 60 percent of the revenues, all Symbolics (Cambridge, MA) had to to do was hold its own against Xerox, LISP Machines Inc., and, most recently, TI. Hardly a prepossessing bunch of competitors, in this market at least.

But now the time is approaching for Symbolics to exploit the market it helped to create. Over the next few years, the installed base of development systems seeded by Symbolics will spawn applications. Lots of new competitors are rushing to provide vehicles for the delivery of those applications. They include, on the one hand, Lucid, Composition Systems, Inc. (page 3), and Silogic (page 14), and on the other, the likes of Sun, Apollo, DEC, Hewlett-Packard, and the great Market Legitimizer itself, IBM. Lucid, CSI and Silogic are developing software to run "AI" on standard machines, while the hardware companies are providing those machines, embellished with AI languages and tools. While not competitors themselves, software companies such as Inference and IntelliCorp are inviting customers to develop on LISP machines and deliver on conventional systems, using software environments to ease the cross-machine transfer. Such a state of affairs would leave Symbolics with the prestige and its competitors with the business.

Symbolics' task is to persuade the world that it offers unique capabilities that are worth paying for, and to provide certain of those capabilities in a system priced for the delivery market. **Result:** Symbolics last week announced a major new release (7) of its software, along with the 3610 Applications Engine, priced from $45,000 down to $31,500 in VAR quantities (100 or more). That still sounds high, but it includes networking, 4MB of main memory and 190MB of mass storage, and superb specs for power, size, noise and reliability. Much of this was achieved by using custom VLSI, reducing the system's component count from 1600 to 600, and boards from three to one.

At the same time Symbolics must price carefully to generate revenues on the high end without driving its customers to the low end. Only upgrade limits, a license agreement and a few missing functions differentiate the delivery vehicle from the development systems. **Result:** Symbolics has sharply cut prices on its development systems, with the low end of the line, the 3640-E, down 25 percent from $65,400 to $49,900. A development system in the same configuration as the 3610AE, the 3640, costs $57,900, down from $73,400.

The new software release runs on both the traditional Symbolics development systems and (in de-fanged version) on the new 3610AE. It has been tuned in lots of little ways noticeable to the user both in increased speed and ease of programming. (Some enhancements noted here are part of release 6, but most existing customers have waited for release 7 to upgrade.) On the speed side, garbage collection* is carried out incrementally rather than in bursts that tie up the machine for minutes at a time. Aside from avoiding this visible downtime, incremental or "ephemeral" GC also improves performance

**"Garbage" is data left over after it's used to derive conclusions; for example, it's the results of all the rules fired to determine the truth of a certain assertion when all you need to know is yes or no. If you later ask the system how it got there, it can easily enough reconstruct the reasoning rather than call it from memory.*

because garbage is collected before it's swapped out to memory, thus reducing unnecessary storing and reloading of garbage. The system uses a certain amount of intelligence to identify garbage and catch it before it's stored. Other less visible tweaks boost overall performance by about 25 percent.

On the programming-tool side, the new software release offers a much richer set of tools both for data- and knowledge-manipulation and for interface-building. Routines for building menus, dialogue boxes and the like are now canned and easily incorporated into a developer's application. Other new features include software diagnostics (looking for bottlenecks in code so that the developer can fix the weakest parts first), table management (so that data organization is optimized, without the user's involvement), mouse-driven and source-level debugging, and file protection (i.e., security).

Staking out the turf

Will Symbolics grab share of the new delivery market opening up? Or will it besmirch its image and get into a pricing war with its low-end competitors? Even though it stands to lose some cachet by going downmarket, its argument of software compatibility up and down the line carries much weight. But now Symbolics must compete with the tools companies (who also help to *sell* its systems) in offering an increasingly rich development environment, complete with Flavors (object-oriented programming), window management, and the like.

Symbolics may feel squeezed, but so do the tools vendors. That's why IntelliCorp is offering SimKit, with other such templates to follow. LISP Machines is making noisy note of its expertise in real-time applications, and its connectibility to the real world of UNIX. Teknowledge and Carnegie Group are converting their tools into C for closer ties to the real world.

Of course, Symbolics' 3610AE announcement won't hold off the other vendors, but it will enable Symbolics at least to share in the delivery market, which should grow far faster than the development market in which Symbolics will remain the leader. As that leader, Symbolics will retain an advantage in influencing the form of the delivery vehicle -- not a mean edge to have.

Looking outside the windows

The infrastructure of Symbolics' announcements comprises several recent moves to position the company more squarely in the real world. It has announced an interface to IBM's SNA (Systems Network Architecture), a support program for value-added resellers, and a network hook-up to support groups of PCs running Gold Hill's GCLISP.

The SNA connection is a promising start, but so far it simply links machines, not applications, together. It doesn't really help a Symbolics application address, say, an IMS data base. The 3610AE somewhat undercuts the GC-LISP link announcement, but for the moment Symbolics is positioning the Applications Engine as a delivery machine, complete with user interface, and the GCLISP-equipped PCs as "naked" machines better suited for engineers than for end-user applications. The VAR program has just started; the first VAR is ICAD, which sells intelligent CAD systems.

All these are small steps, but they send a clear message: Symbolics wants to interconnect to and serve environments outside its own.

COMBING SPAGHETTI WITH RECODER

New software is exciting to technologists, programmers, and journalists. But to a true salesperson, an installed-base market is far more exciting, by virtue of sheer size. Consider the 50 billion lines of COBOL, tended by millions of software engineers, that run much of this country. Some people, of course, think they're strangling the country. Most code is reasonably straightforward when it's first written, but as it picks up patches and paper clips over the years it deteriorates into a frightening mess.

One of the companies dealing with this problem, Language Technology Inc. of Salem, MA, takes the problem down to its essentials. Its Recoder parses a COBOL program, graphs it to determine its structure, and recodes it into structured COBOL. While Recoder doesn't "understand" the program, it does "understand" its structure, and produces a new program that's guaranteed to do what the old one did, bugs and all.

There are numerous beneficial effects. First of all, the resulting code is far easier to maintain. Think of structured code as something like text in an outline processor (or auditing functions in a spreadsheet): You can determine the dominance and "reach" of any given statement or set of statements. That is: If I change this little thing here, where will it have an impact aside from where I intend it to? This is perhaps the most important question that can ever be asked in maintenance, but with unstructured code, you can't tell.

The mean time to failure after a major patch of an aging COBOL program, LTI's president Bill Engel tells us, is normally 90 minutes. With "Recoded" COBOL, it's a matter of days. The programmer can find the piece he wants to fix, and then be sure that he is "fixing" only what he intends to.

Structured code is also easier to understand, document, and find bugs in. Although Recoder won't take bugs out (it delivers back to you exactly what you gave it, albeit in structured form), it will point them out or make them easier to spot. Additionally, structure makes it easier to figure if the system is doing what you want it to, or what you *think* it's doing. That go-to that goes nowhere, or those three little lines that alter the programmer's salary, will reveal themselves incontrovertibly.

In fact, the idea makes so much sense that IBM in January launched a similar product, COBOL Structuring Facility. However, at this point it's unlikely IBM's product will do much more than legitimize the market, as it works only with IBM's new VS COBOL II, or a minuscule though growing portion of the installed base. Moreover, COBOL/SF is not fully automated; it works best with substantial pre-processing (like a pre-wash cycle?).

Energetic Engel

Recently arrived from Xerox Computing Services, where he was marketing and sales vp, Engel has taken an interesting idea and put it on a commercial footing. The company, founded by Harvard philosophy graduate Eric Bush, 34, had been selling Recoder since early last year as a service, "Ten to 50 cents a line to wash your COBOL." For a typical insurance company, with 5 million lines of code, that worked out to half a million. (Nice work if you can get it.) Besides, who wants to hand over code to an outside service? Response was lukewarm at best.

And finally, while management liked the idea, it was the operating departments who made the purchase decision, and they, for reasons of turf, were less enthusiastic. Now, as a package costing $150,000 for a three-year license, Recoder is sold to management, which presents it to users to take in-house and run under their own control, at their own pace. (COBOL/SF sells for $125,000, or $12,500 monthly.) Typically, it takes a couple of minutes to restructure a normal-size program segment -- a few thousand lines of code -- and an hour or two to test it thoroughly thereafter (generally by running the old package and the new improved version against the same tasks, and then comparing the results.) Since the recoding process is totally automated, except for selecting the code to be altered and specifying a few parameters on a menu (built with IBM's ISPF dialogue manager), the job can be run late at night on deadtime.

As it happens, even with structured code, maintenance is generally a horrifying mess of go-tos and other coding no-nos: Payroll won't wait for an elegant solution. In practice, Engel expects, Recoder will be run against working software on a continuing basis, rather than once and for all.

Does the restructured software actually run better? Not necessarily. Recoder uses cunning proprietary "AI" techniques to determine and graph the code's structure, and then writes it out cleanly. But it doesn't optimize, as a (skilled) human programmer might. In fact, optimization may be less than straightforward -- and downright cryptic to anyone who didn't write it himself. The minor penalty in performance is insignificant compared to the improvement in repair time, uptime, and documentability.

Resisting Recoder

Bush and Engel want to get a few reference accounts first, and worry about smaller companies later. The company's eight-person salesforce is calling directly on 2700 mainframe sites from a Computer Intelligence listing. Early customers include Federal Express -- only 10 years old, but with an ample supply of spaghetti code already -- Phillips Petroleum, TRW, Merrill Lynch, GTE, and the Mellon Bank. (Early investors -- for $4.5 million -- include Citicorp Venture, Inco Venture, Battery Ventures, and Continental Illinois.)

Even though Recoder does all that's promised, it's not an easy sell. One customer we checked with reports that his engineers complain that Recoder produces "overstructured" code. That is, they say, they had reasons for those go-tos, and they find Recoder mildly insulting. What about restructuring old code written by former employees? "Well, that's a mess, sure. But all you get with Recoder is a *structured* mess."

Everyone likes productivity, yes, but it's supposed to come from human effort, not from some machine. Right?

RELEASE -1.0

Q. What is French? A. Latin 2.5.

RELEASE 2.0: SILOGIC RESTRUCTURED

We last wrote about expert-system tool vendor Silogic in September 1984, when we noted carefully that the company was not yet shipping. Last summer, says Al Barkovsky, then marketing vp and now president, "I had to recommend to the board whether to shut the whole thing down." He brought in a consultant from Arthur D. Little who vouched for the soundness of the technology the company was developing, and so Barkovsky asked the board for control of the company and another $1 million, on top of $2.6 million already sunk.

Then, he says, "I stopped all development and told them to turn it into a product." That's what's going on now; more on the product below. There remained one other big problem: Silogic has the visibility of a submerged hippopotamus. Funded by two local Los Angeles real-estate tycoons, the company totally lacks the roster of Stanford/MIT professors, venture capitalists and "strategic investors" that have proved so helpful to the established AI companies who would not even deign to consider Silogic a competitor. No problem, but potential customers and resellers likewise don't consider Silogic a serious player. For the last few months, Barkovsky has been touring the country for four packages of prospective-board-member-plus-customer-plus-money, as well as one brand-name venture capitalist who need not be a customer. "We aren't very original," says Barkovsky; "we want a bank, a consumer goods company, an aerospace company, and an auto maker."

For the moment, Barkovksy is only in "talks" and can't name any names, but he has a number of beta customers, at least. One is Latham Process Control of New York, which is building a production control system for printing companies; others are aerospace/defense units at Eaton Corp. and the Jet Propulsion Lab in Pasadena. Most interesting, both commercially and politically, is Sumitomo Electric Industries of Japan. Sumitomo, acting as a distributor, has 15 copies of the product, mostly in use for diagnostic applications. Although Silogic won't confirm it, we understand that one of those copies has found its way to ICOT, the vaunted Japanese Fifth-Generation project, so far more successful in hardware than in software.

Knowledge Workbench

Silogic's product is the Knowledge Workbench, a suite comprising a Natural-Language Processor (NLP), the Inexact Reasoning Module (IRM, a Prolog-based backward-chaining expert system engine), and the Knowledge Base Management System (KBMS). All three modules together sell for a pricy $60,000, support included; VAR discounts start with the second unit. KBMS gives the Knowledge Workbench its unique and attractive flavor. It's a relational-dbms-like knowledge base built in Prolog, and looks at standard SQL-type relational data bases *as if* they consisted of Prolog clauses. That produces an unusual closeness between the data and the inference engine, and makes it extremely easy to reason about large volumes of data stored in "standard" form. (It's not so easy to get at "old-standard" data bases, such as IMS or IDMS, or ISAM and VSAM files.)

For example, suppose you have a hierarchy: Juan reports to Alice, Alice reports to Fred, Fred reports to Dave, and Dave reports to his mother. Each of those relationships can be represented as a single Prolog clause. Then add the rule that if A reports to B and B reports to C, then A reports to C. (You could also define "reports-direct," where this would not be the case.)

This would enable you to determine that Juan ultimately reports to Dave's mother. It's not something you could easily accomplish with a simple relational data base, or with a collection of rules, which is just not the best way to represent data that fits well into tables. You need the combination.

It also so happens that Prolog is well suited for the job, since Prolog assertions are structured almost like data entries in a relational table:

```
reports_to (Alice, Juan).
reports_to (Fred, Alice).
reports_to (Dave, Fred).
reports_to (mother, Dave).
```

would do the trick here.

Knowledge Workbench's natural-language interface handles not just standard data queries and assertions, but can also be used to build rules. The good news is that it deals well with data and rules; the bad news is that it still seems to have little "interest" in talking to other applications.

In other words, the Knowledge Workbench is driven by interaction with a user rather than triggered by interactions with other software. To be sure, no canned commercial tool really works that way at this point, with the possible exception of mdbs's Guru (Release 1.0, 32 December 1985). It's the same problem we noted earlier with electronic publishing systems: They talk well to users, but the component applications rarely talk to each other.

Backwards only

Overall, Knowledge Workbench seems to be a fairly standard backward-chaining inference engine, best suited for rule-driven problems and similar in that respect to Teknowledge's product line (Release 1.0, 14 March). However, Silogic claims, it has done a better job of handling negation and checking for logical consistency. (Just wait for everyone's next release.) Like Teknowledge, Silogic is attempting to deal with real-world problems in real-world environments. Teknowledge has written its systems in C to run on standard hardware. Silogic is still using Prolog, but it delivers its product on many 68000-based systems such as Sun, Masscomp, Plexus, and NCR; the choice of others will be due more to marketing than technical considerations. The data base interface currently is tuned to work with the Unify relational dbms, but there again the addition of other SQL complements would be easy. Less easy would be porting the system to IBM mainframe equipment, (high-end) PCs, or the RT PC. But Barkovsky is at least thinking about it.

Late into an increasingly competitive market, Barkovsky hopes to start commercial shipments this fall, primarily to VARs (a philosophy dating back to our first encounter with the company). Silogic will supply the systems expertise; the VARs will provide applications expertise and presence in distribution channels. Silogic is also finding a market among hardware vendors who want to put AI in their catalogues, although Barkovsky has his doubts, as do we, whether their interest extends beyond an earnest wish to put a check in the "AI" column in their brochures.

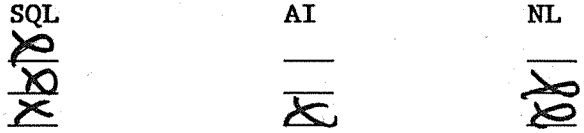

RELEASE 0.5

SMARTNOTES *(Remember to put title in boldface.)*

SmartNotes, from Personics of Concord, MA, *(remember company reference)* is a more general approach to the problem that Turner Hall's Note-It solves for Lotus spreadsheets in particular -- with the broader market and the lesser functionality that particularity implies. A stroke of marketing genius, they are *(it is?)* the electronic equivalent of 3M's Post-It *(does this need a TM sign?)* notes; they stick to anything -- directory listings, spreadsheet cells, calendar dates, names in a text file, what have you. SmartNotes was developed by two professors, one at Carnegie-Mellon and one at Princeton, who were having trouble communicating electronically as they attempted to write and edit joint papers on algorithm development. Typing notes into the text was intrusive. *(...as shown here)*

(How do they work?) SmartNotes is an excellent example of the interstitial programming mentioned on page 9, *(make sure page reference is correct)* as are most terminate-and-stay-resident packages. Each note is attached not to a "physical" place in a file but rather to a character string or *(treatment of cells and formulas?)* a spreadsheet cell address. In text mode, the Note follows the string wherever it moves in the file. Whenever the string shows up on the screen (and the notes are turned "on"), the SmartNote pops up. In spreadsheet mode, SmartNotes can still read only what's *on* the screen, so it won't notice formulas, but it can read column and row headings to determine which cells are on display. *(What about spreadsheet restructuring?)* It attaches to a particular cell, and adjusts whenever rows or columns are added or deleted, as does Note-It.

The biggest problem: What happens if a string used as an anchor is edited or recalculated? *(Remember to note price, memory requirements.)* The $79.95, 90K program will warn the user that he's about to lose his anchor. Another feature -- or a bug -- is that the Note will appear in reference to, say, "Alice Haynes," "1986 revenues," or "$149.95" every time that string appears on your screen. The precise string referred to is under the control of the user.

Just imagine having a pink box pop up, pointing to a file called "co.let" in your directory, explaining that this is the source file for a mail-merge of letters addressed to the names in "co.names." *(We could have used that a couple of months ago.)* On an agenda you send to your secretary, *(copies?)* who must have her own copy of the program, May 17 could be marked: "Change this flight to Tuesday. Move up reservations at Rickeys." Or one you send to a colleague: "Do you think this is in poor taste?" with the response, "Yes -- but he deserves it." *(Is this too many cute examples?)* As people start using e-mail, SmartNotes should come into their own -- and they don't fall off as easily as those paper ones. *(Put a snappy close here.)*

SMART MAPS

For most people, working with maps has been akin to drawing graphs with a "draw" package: Figure out your data, then draw the graph and extend the lines or bars to the right length. No one does that anymore, of course. But they still download map outlines from libraries of shapes, and color or shade the areas within the maps one by one to illustrate their data.

While many vendors offer picture files -- maps, icons, stick figures, chart formats and the like -- to be added onto data-based graphs, these images have no relation to any data. For technically minded people, Strategic Locations Planning of San Jose has long offered Atlas, a data-rich mapping program tailored for market analysts, an outgrowth of the firm's business consulting to retailers and bankers expanding to new locations. Strategic Locations' people come from the Boston Consulting Group and sell their product mostly through ads in American Demographics and other such media.

Now the broad commercial market is getting its own "friendly" implementation of the idea. DecisionResources of Westport, CT, a longtime vendor of pc graphics packages (Chart-Master, etc.), has produced an appropriate integration of map files with, well, smarts. Like Atlas's, its Map-Master maps are linked to underlying data -- a variety of demographic and marketing statistics, or your own proprietary data, by town, state, census area, ZIP code, or whatever. While it lacks some of the more esoteric features of Atlas (dot-density maps, for example), DR's Map-Master is easy to use for the average joe making a presentation.

Simply fill in your own sales figures in the form linked to the areas you cover, or load in an appropriately keyed 1-2-3 file, and Map-Master will draw you an appropriately colored map showing where your salesforce is above quota, and where below. Tell it the number of stores you have in each area, ask it to divide by the local population of women 18-49, and you'll be able to see where there's need for another outlet of your Slinky Softwear stores.

Map-Master costs $395 including basic data files, with extra files for $95 apiece. Atlas costs $449 in a similar configuration. A third product, Randmap from a Rand McNally subsidiary in Chicago, costs $995.

All three products, like a graphics package or a project manager, use graphics as more than mere decoration. They generate data-driven maps, just as more familiar graphics packages produce data-driven bars, lines, and pies. They do what analytical/presentation software is *supposed* to do; they turn data into communicable information.

SQUEEZE ME!

From the folks who brought you Note-It (Turner Hall, a catalogue house in Cupertino owned by Symantec), comes SQZ! A data compressor that the company has priced low enough -- $79.95 -- to avoid encouraging competition, SQZ! looks like a no-brainer purchase for anyone with 1-2-3 or Symphony spreadsheets. Since it stays in memory and works automatically any time a file is loaded or stored, SQZ! actually reduces disk access time, as well as the amount of data that needs to be stored or transmitted, saving disk space or telecommunications costs. Files are compressed down to 5 to 20 percent of their original size.

The only problem with SQZ! is that it renders files invisible in a Lotus directory: If SQZ! is on, the Lotus menu line will list only Squozen files; if SQZ! is off, it will list only unSquozen files. Until you're used to it, you may think you've lost some files. The reason for this is that SQZ! captures Lotus's loading and storing routines, and converts the Lotus WKS or WK1 suffix to WK$ or WK!

iLINK ME -- BUT NOT TOO TIGHTLY!

Lotus recently announced the demise of its vaguely defined joint venture with Cullinet, which would have enabled Lotus programs to trade data (not just files) easily with Cullinet's widely used IDMS and IDMS/R data bases. It reminds us of a similar venture between Apple and Cullinet several years ago (involving the Lisa!): Each company thought the other should be grateful and do all the work, and in the end no one did anything.

Instead, Lotus recently bought rights to iLINK, another micro-mainfame link product, from InfoCenter Software of New Paltz, NY. iLINK connects only to a bunch of information-center-style data bases -- RAMIS, Nomad, Focus, etc. -- most of which are already designed to be "friendly" and less in need of auxiliary links to end-users in the first place. The world of IMS, IDMS and the like, the repository of the vast majority of corporate data, remains unreachable directly by iLINK. But perhaps the *real* customers -- the dp managers, not the end-users -- like it that way.

The most general solution we're aware of comes from Answer Systems, formerly Informatics' Software Products Group and now part of Sterling Software. Answer, based in Canoga Park, CA, offers Answer/DB and associated PC products, which let users of 1-2-3, dBASE and other PC spftware import data directly from IMS, DL/1, and VSAM data bases -- IBM models which account for about 65 percent of the mainframe dbms market -- and from Cullinet's IDMS, which accounts for much of the rest. (The drawback: Answer/DB costs $20,000 to $30,000, plus $550 for each pc file format. iLINK costs about the same per mainframe, but without per-cpu charges.)

Yet mainframe dp managers aren't all that eager to have end-users mucking about in their production data bases, allowing some junior MBA to halt payroll to do a staffing analysis. Those mainframe managers are the ones who will be making the purchase decision for Answer/DB, especially since a recent (non-exclusive) marketing agreement between Answer and IBM. They may question Answer president David Saykally's contention that his product line fully addresses the most pressing issues of data security and priority.

Lotus, with iLINK, takes a far less disruptive approach, simply giving users access to a nicely segregated information center data base -- which generally has tightly controlled access to a company's production data. Maybe this is Lotus's way of endearing itself to dp managers... In essence, Lotus is qualifying its customers by de facto requiring that they have gone to the effort of offloading corporate data to an information center. "The non-information-center market is second-priority," says Don McLagan, vp and general manager of Lotus's Information Services business. Apart from the technical issues, "there's so much administrative work to make the data accessible that we prefer to target the 6000-odd sites who've already done it."

SUN, APOLLO TIE UP WITH TI. WHY?

Sun Micro and Apollo have both recently announced marketing agreements with TI. TI will use Sun's rapidly spreading standard, the Network File System, to link its Explorer workstations. In a somewhat more intimate relationship, TI will sell the Explorer to Apollo for integration with that company's Domain line of networked workstations. With third-party LISP vendors

such as Lucid around to turn 68000-based workstations into de facto LISP machines, you'd think these companies would be competing with TI rather than cooperating with it.

While Sun is adamant in favor of standard architectures and is licensing NFS to all comers, Apollo at least may want to be first in line when TI starts making its LISP-on-a-chip available to third parties. And of course some people think LISP *is* a standard. Certainly a low-end LISP-on-a-chip, suitable for embedding anywhere, would help make it so.

PROGRAMMED PR

One of the more visible vendors at the Microsoft CD-ROM conference last month was Computer Access Corp. of Belmont, MA, suppliers of BlueFish, a text-search program (which it likes to call a "publisher's tool") that works just as well with regular old disks as with CD-ROM. Its demo offered the book *The New Papyrus* on CD-ROM, showing how you could search the text almost instantaneously. Far more interesting to our audience -- but not shown because it's on plain old floppy disks instead of CD-ROM -- is the Artemis data base, from Artemis Media Corp. of Framingham, MA. Formed by former Prime Computer marketing manager Steve Seltzer, Artemis is building a data base comprised of editorial calendars solicited directly from publications, enhanced with editors' names, closing dates, phone numbers and other data about each periodical. (aRTeMiS refers to Real-Time Media Search and also to Artemis, the Greek goddess who (mistakenly) shot her lover Orion, obscured beneath ocean waves, with fatal accuracy.)

Suppose you've just joined Expert Evaders, an AI tax-software firm, and your first job is to promote the ceo, a former CPA and income-tax whiz. You search for all the calendars including both "tax*" (tax, taxes, taxation, etc.) and "software." Then you call the appropriate editors... Or with a limited budget, you might want to run just three ads around the time the product you have just announced will be shipping, and to place them near appropriate editorial content. You can instruct the system to find only articles running six months from now.

Given Seltzer's background in the computer industry, the computer books are among the first periodicals to be compiled, with others to follow. A subset -- such as that for the computer industry -- costs $995 including search software; the entire data base, covering 20,000 distinct issues of periodicals, costs $19,000. Mundane, yes, but it's easy to see a lucrative market for this service. Revenues will be recurring, as the data base needs frequent updating -- at $5000 per year. Initial clients include ad agency Foote, Cone & Belding, pr firm Newsome & co., and, of course, Seltzer's old employer Prime. *Now who's running an article on industry newsletters?*

IntelliSHUFFLE

IntelliCorp recently made some common knowledge tangible when it made president and ceo Gene Kromer president and coo, and evp Tom Kehler chairman and ceo. Kehler, a young (39), likable and not terribly organized former physicist, is far better-suited to lead than to manage. By contrast, Kromer, 52, a comfortable, seen-it-all veteran of 20 years at Texas Instruments be-

fore he joined IntelliCorp in 1984, is far better-suited to manage than to dash around the country impressing investors and gleaning contacts, as Kehler will now be free to do. Every relationship has its tensions known only to those involved, but this trade in positions is about as straightforward as such an event can be. Whatever tensions exist relate to the retirement of co-founder Laurence Kedes as chairman of the board (although he remains a director). Kedes, a professor of medicine at Stanford, was most closely linked to the genetic-engineering side of the company (the whole company was formerly called IntelliGenetics), part of which was recently sold to Amoco.

XEROX AI: CINDERELLA TO THE BALL?

Sometimes, things have to get worse before they get better. Such seems to be the case at Xerox's Artificial Intelligence Strategic Business Unit (or AIS), long the ugly stepsister to the company's glamorous (if avowedly uncommercial) Palo Alto Research Center. The AI SBU stuck determinedly to its own, non-standard InterLisp, and seemed incapable of making headway in a now rapidly growing market. Although Xerox paid lip service to the role of AI and proudly talked of the achievements of PARC, it seemed unwilling to devote adequate resources to the commercialization of it through AIS.

This month things came to a head, with the departure of the third significant person from the group. First, manager of product development Beau Sheil left his post for a less defined future as a consultant with Sutherland, Sproull & Associates, Menlo Park. The decision was his own, but we suspect it reflects mutual unhappiness. Another departure was that of Steve Purcell, chief engineer of the LISP implementation for the 1186, Xerox's AI delivery system. Purcell, an outside consultant with a non-Xerox-standard contract that somehow slipped through for years, simply couldn't get it renewed for next year. And finally, SBU manager and acting sales chief Gary Moskovitz, formerly with Mattel, left last week for a better oppportunity as president of Bell & Howell's computer-output microfilm business.

Clearly, something had to be done. Last weekend, Xerox top management held an out-of-the-ordinary planning meeting to ratify what had been under discussion for months. The outcome: Two seasoned Xerox managers, of a caliber likely to requite continued top-management support, are coming in to run the SBU. Acting general manager Joe Backer, formerly Engineering and Scientific business and program planning manager at Xerox and previously with IBM and SDA, a CAD/CAM startup, has both the management skills and the technical grounding to earn the respect of his new unit. New sales manager Gary Aslin, an ex-Navy fighter pilot, has the energy to earn the respect of customers. The unit's overall development staff will increase by at least 70 percent by year-end, and will work especially closely with the technically fertile PARC scientists, pending the arrival of a new development manager.

All in all, Cinderella's future is not yet assured, but at least she'll get a chance to show her mettle.

PHONE NUMBERS: see page 23

RELEASE 1.0 CALENDAR

Sometime this week... **"Hackers"** - on your local public TV station. Documentary filmed at the Hackers Conference last summer. For times and stations, call (415) 751-8808.)

April 21-24 — **Sixth International Conference on Decision Support Systems** - Washington, DC. Sponsored by the Institute of Management Science. Call Julie Eldridge, (401) 274-2525.

April 28-30 — **Avignon 86** - Avignon, France. Sixth international workshop on expert systems and their applications. Contact: Jean-Claude Rault, Agence de l'Informatique, Tour Fiat - Cedex 16, 92084 Paris-La Defense, France; telephone: (011 331) 47 96 43 14.

April 28-May 1 — **Comdex Spring** - Atlanta. Contact: Linda Yogel at the Interface Group, (800) 325-3330.

April 29-May 1 — **AI '86** - Long Beach, CA. Contact: Norma Ashman at Tower Conference Management, (312) 668-8100.

May 4-7 — **Fourth National Conference on EDP Software Maintenance** - San Francisco. Sponsored by the Data Processing Management Association. Contact: Conference Manager, US Professional Development Institute, (301) 445-4400.

May 5-6 — **AI Conference** - Stamford, CT. Organized by K. C. Branscomb, (415) 851-4735, in conjunction with the Gartner Group. For registration information call Ashley Pearce at the Gartner Group, (203) 967-6757.

May 5-7 — **Federal DP & Communications Conference** - Washington, DC. Sponsored by the Interface Group. Contact: Rachel Domba, (617) 449-6600.

May 6-8 — **National Online Meeting** - New York City. Contact: Joanne Loreti at Learned Information, (609) 654-6266.

May 6-8 — **Videotex '86** - Dallas. Contact: Regina Kreischer at Online Conferences, (212) 279-8890.

May 12-13 — **Personal Computing Conference** - Nashville. Contact: Ashley Pearce at the Gartner Group, (203) 967-6757.

May 12-14 — **OpticalStorage Forum** - Denver, CO. Cosponsored by DISK/TREND and Freeman Associates. Contact: Darlene Plamondon at Catlidge & Associates, (408) 554-6644.

May 14-16 — **Info Show** - Los Angeles. Contact: R. Waitzman at Cahners Exposition Group, (203) 964-0000.

May 18-21 — **Hammer Forum** - San Francisco. "Shaping the Second Revolution in End-User Computing." Call Pamela Davis, (617) 354-5555.

May 19-21	**IIA Midyear Management Conference** - Denver. "Information Everywhere..." Call Virginia Nelson at the Information Industry Association, (202) 639-8262.
May 20-21	**PC Products 86: Tech Trends** - Dallas. By Future Computing. Contact: Charlene Stevenson at (214) 437-2400.
May 21	**National Data Corporation presentation for analysts** - Atlanta. Call L. C. Whitney's office at the company, (404) 329-8506.
June 1-4	**Summer Consumer Electronics Show** - Chicago. Contact: Electronic Industries Association/Consumer Electronics Group, (202) 457-8700.
June 1-4	**AICPA Annual Microcomputer Conference & Exhibition** - Boston. Call Linda Franco at the American Institute of Certified Public Accountants, (212) 575-6451.
June 8-10	**1986 Artifical Intelligence Forum** - Stamford, CT. Sponsored by New Science Associates. Contact: Bruce Barlag, (203) 855-0101.
June 8-11	**World Computing Services Industry Congress** - Toronto. Call Phyllis Cockerham at ADAPSO, (703) 522-5055.
June 9-11	**Third Annual Software Industry Conference** - Monterey, CA. Sponsored by the Gartner Group. Contact: Ashley Pearce at (203) 967-6757.
June 10-12	**Comdex International** - Nice, France. Contact: Pete Young at Interface Group, (800) 325-3330, or Charles Vervoord, Amsterdam, (31) 20-6219 41.
June 16-19	**National Computer Conference** - Las Vegas. Sponsored by AFIPS. Contact: (800) NCC-1986.
July 9-11	**PC Expo** - New York City. Contact: Steve Gross at PC Expo, (800) 922-0324 or (201) 569-8542.
July 21-25	**European Conference on Artificial Intelligence '86** - Brighton, UK. Contact: Luc Steels, Vrije Universiteit Brussel, AI Lab, Pleinlaan 2, 1050 Brussels, Belgium; telephone: (011 320) 2-641 2971.
July 22-24	**Future of Optical Memory Technology** - San Francisco. Sponsored by Rothchild Consultants. Call Judith Hanson at (415) 626-1133.
August 4-6	**IIA Conference** - Cambridge, MA. "Beyond databases: Knowledge delivery systems." Call Linda Cunningham at the Information Industry Association, (202) 639-8260.
August 11-15	**AAAI-86** - Philadelphia. Call Claudia Mazzetti, American Association for Artificial Intelligence, (415) 328-3123.

August 19-22	**International Conference on Parallel Processing** - St. Charles, IL. Sponsored by IEEE and Penn State. Contact: Earl Swartzlander, TRW, (213) 535-4177.
September 2-5	**Comdex/Australia** - Sydney. Yet another Comdex. Call Linda Yogel at the Interface Group, (800) 325-3330.
September 7-10	**ABCD Breakway '86** - Lake Buena Vista (Orlando), FL. The third annual convention of the Association of Better Computer Dealers. Contact: Terri Rojas at (312) 693-ABCD.
September 8-10	**NCC Telecommunications '86** - Philadelphia. Sponsored by AFIPS. Contact: (800) NCC-1986.
October 20-22	**UNIX Expo** - New York City. Contact: Don Berey or Bob Birkfield at National Exposition Co., (212) 391-9111.
October 27-29	**DPMA International Computer Conference** - Atlanta. Sponsored by the Data Processing Management Association. Contact: Doris Dodge or Marion Okon, at (312) 825-8124.
November 2-5	**ADAPSO 65th Management Conference** - Phoenix. Call Phyllis Cockerham at ADAPSO, (703) 522-5055.
November 10-14	**Comdex Fall** - Las Vegas. Contact: Linda Yogel, Interface Group, (617) 449-6600.
February 22-25 (1987)	**Tenth Annual Personal Computer Forum** - Phoenix. For information, call Sylvia Franklin at EDventure Holdings/ Release 1.0, (212) 503-5517.

Please let us know of any other significant events we should include.
-- Laurin Henchey

PHONE NUMBERS

David Saykally, Answer Systems, (818) 716-1616
Steve Seltzer, Artemis Media, (617) 872-0722
Dave Norman, Businessland, (408) 554-9300
Mike Stock, Composition Systems, (914) 592-3600
Jamie Rapperport, DecisionResources, (203) 222-1974
Tom Kehler/Gen Kromer, IntelliCorp, (415) 965-5500
Bill Engel/Eric Bush, Language Technology, (617) 741-1507 or (800) RECODER
Jerry Kaplan/Don McLagan, Lotus Development, (617) 577-8500
Marc Peterson, Personics, (617) 897-1575 or (800) 445-3311
Ann Veit, Randmap, (312) 673-MAPS
Al Barkovsky, Silogic, (213) 337-7477
Stephen Poizner, Strategic Locations Planning, (408) 985-7400
Scott McNealy, Sun Microsystems, (415) 960-1300
Beau Sheil, Sutherland, Sproull, (415) 326-7438
Bill Hoffman, Symbolics, (617) 259-3600
Tom Byers, Turner Hall (Symantec), (408) 253-9607

SUBSCRIPTION FORM

Please enter my subscription to **Release 1.0** at the rate of $395 per year in the U.S. and Canada. Overseas subscriptions are $475, airmail postage included. Payment must be enclosed. Multiple-copy rates on request.

Name_____

Title_____

Company_____

Address_____

City_____ State_____ Zip_____

Telephone_____

SC

Please fill in the information above and send with your payment to:

Sylvia Franklin
Release 1.0
Ziff-Davis Publishing Company
One Park Avenue
3rd Floor
New York, NY 10016

If you have any questions, please call us at (212) 503-5517.